**Article:**

Carneiro, Pedro, Lopez Garcia, Italo, Salvanes, Kjell et al. (1 more author) (2021) Intergenerational Mobility and the Timing of Parental Income. Journal of Political Economy. 712443. ISSN: 1537-534X

https://doi.org/10.1086/712443

# Intergenerational Mobility and the Timing of Parental Income

Pedro Carneiro

University College London, HCEO, IFS, CEMMAP, FAIR-NHH

p.carneiro@ucl.ac.uk

Italo Lopez Garcia

RAND Corporation, IZA

ilopezga@rand.org

Kjell G. Salvanes

Norwegian School of Economics, HCEO, IZA

kjell.salvanes@nhh.no

Emma Tominey

University of York, HCEO, IZA

emma.tominey@york.ac.uk

June 2020

Abstract

We extend the standard intergenerational mobility literature by examining the relationship between adult outcomes of children and the timing of parental income during their childhood years, using data from Norway. We find firstly that, conditional on permanent household income, the child's human capital is higher in households where income is balanced between the early childhood and late childhood years than in households with a more imbalanced income profile. Second, compared to the early and late period of childhood, income in the middle period has a relatively low productivity.

* We thank participants in several seminars for useful comments. Carneiro gratefully acknowledges the financial support from the Economic and Social Research Council for the ESRC Centre for Microdata Methods and Practice through grant ES/P008909/1, and the support of the European Research Council through grant ERC-2015-CoG-682349. Salvanes thanks the Research Council of Norway for financial support. Tominey acknowledges the Research Council for Norway, the Leibniz Association, Bonn, in the research network 'Non-Cognitive Skills:Acquisition and Economic Consequences' and the British Academy for support.

# 1 Introduction

There is a large empirical literature examining the intergenerational transmission of economic status (for recent surveys see Solon (1999); Black and Devereux (2011); Björklund and Salvanes (2011)). It is possible to find estimates of intergenerational mobility for various outcomes, and for virtually every country in the world where data linking parents and children is available. Most estimates come from simple models relating a measure of child income and a measure of parental income, such as:

$$y_c = \alpha + \beta y_p + u, \tag{1}$$

where $y_c$ is a measure of the child's income, $y_p$ is a measure of parental income, and $u$ is a residual.

Standard economic models of intergenerational transmission justify the use of equation (1) (e.g., Becker and Tomes (1979), Becker and Tomes (1986)), but they collapse the childhood years to a single period of life. More realistic models of parental investments in children distinguish several stages of childhood (e.g., Cunha and Heckman (2007), Cunha, Heckman, and Schennach (2010), Cunha (2013), Caucutt and Lochner (2012)). They make clear the idea that the whole history (in particular, the timing) of parental investments in children may be as or more important than the total amount invested during the childhood years. In other words, the timing of income may be as or more important than a single measure of income, and the model of equation (1) is missing this issue.

This paper extends the literature on intergenerational transmission by examining the relationship between adult outcomes of children and the timing of parental income during their childhood years, using administrative data from Norway. We focus our attention on two child outcomes: years of schooling, and earnings at age 30 (although we also present results for dropout from high school, college attendance, IQ and teen pregnancy). To simplify our empirical procedure, we divide childhood in three periods: early (ages 0-5), middle (ages 6-11) and late (ages 12-17).[1] We find that the timing of income matters, over and above permanent income.

Our empirical procedure compares outcomes of children across families with different income profiles across the three periods of childhood, but with the same permanent income.[2] Although our comparisons are across families, it is easier to describe our results in terms of a single family facing different counterfactual income profiles.

Our measure of permanent income is a statistical one, defined as the sum of (deflated) family

[1] Our results are robust to dividing childhood into more periods.

[2] When we represent our results we keep both fixed both permanent income and income in one of the three periods of childhood, e.g., the late period. Then we can compare child outcomes for children growing in families where, for example, incomes are (relatively) high in early childhood and low in middle childhood, with child outcomes for children experiencing low early childhood income and high middle childhood income. Notice that conditional on permanent income during the childhood years and income in one period of childhood, incomes in the other two periods of childhood are linearly dependent: when one rises the other must fall.

income observed during each of the first 17 years of life of the child, discounted to the time of the child's birth (we also consider other measures that consider the whole lifetime of children and parents).[3] As shown in Cunha and Heckman (2007) and Caucutt and Lochner (2012), with perfect credit markets and no precautionary savings the timing of income should not matter. In contrast, the timing of income should be important if families face restrictions to borrowing or saving. This is the main rationale behind our empirical procedure.

We show that, keeping permanent income and income in late childhood fixed, child's schooling and earnings improve if a family experiences relatively high (counterfactual) income in early childhood and low income in middle childhood, rather than a relatively low income in early childhood and high income in middle childhood. Stating it more simply, child outcomes improve as (counterfactual) family income rises in early childhood and falls in middle childhood. We also show that a child's schooling and earnings increase when her family income falls in the middle years and increases towards the later years of childhood. Compared to the early and late period of childhood, income in the middle period seems to have low productivity.

Finally, children in families with a balanced profile of income between early and late childhood, keeping fixed permanent income, experience better adult outcomes than those in families with income mainly "front-loaded" in the early period, and those in families with income mainly "back-loaded" in the later period of childhood. In other words, we find that there is an inverse U-shaped relationship between income in the early period relative to income in the late period (which means keeping fixed permanent income and income in middle childhood), and the child's education and labour market outcomes in the adult years.

These empirical patterns are consistent with a model with income uncertainty, partial insurance and imperfect credit markets (Blundell, Pistaferri, and Preston (2008), Cunha and Heckman (2007), Caucutt and Lochner (2012)), and complementarity between investments in children across periods (Cunha and Heckman (2007), Cunha, Heckman, and Schennach (2010)). Income uncertainty, together with partial insurance possibilities such as an inability to borrow and save freely (e.g. for precautionary motives), lead to a setting where investments in children react to parental income shocks in each period. Complementarity between investments taking place in different periods means that human capital is maximized when there is a balanced flow of investments.

In particular, this combination of credit market and technological factors can result in a model where a balanced flow of income shocks may lead to higher human capital than an unbalanced history of shocks. This would be consistent with an inverse U-shaped relationship between the education of the child and the amount of income that is front-loaded (or back-loaded) in the initial (late) period of childhood. However, our results also suggest that income in the middle period of childhood may be relatively unproductive, compared to income in earlier and later periods.

[3] In other words, this is not a measure of expected future family income at birth (or in any other year), or even a measure of assets, on which parents base their consumption and investment decisions. It is a purely statistical measure used to construct a test of whether the timing of income matters.

As mentioned above, we expand equation (1), by replacing the single regressor $y_p$ with multiple measures of income, measured at different stages of childhood. In order to do this we need data on the entire family income history for a large number of children, which is rarely available. We use data from Norwegian registries for children born during the 1970s, which allows us to link an individual's outcomes as a young adult to the whole history of parental income during the childhood and adolescent years.

We briefly discuss the extent to which we can interpret the association between the child's education and the timing of parental income as causal. We face three (related) challenges. First, the precipitant of the income shock, not the income shock itself, could be the driver of child outcomes. Second, the timing of income could be a choice variable, potentially correlated with investments in children. For example, there are periods when parents take time off work to take care of their children, such as in the case of maternity leave. Third, there is heterogeneity across parents, which can be simultaneously related with income profiles and investments in children. For example, parents on steep income profiles may be more able than those with flatter income profiles. Our discussion is supported by an extensive analysis of these issues, which is included in an appendix.

There exists already some literature explicitly examining the role of the timing of income in the formation of human capital. Some of these papers use survey data from the US and Germany, and rely on relatively small datasets (Duncan, Yeung, Brooks-gunn, and Smith (1998), Levy and Duncan (2000), Jenkins and Schluter (2002), Carneiro and Heckman (2003), Caucutt and Lochner (2012)). Others use much larger register data for Denmark and Norway (Aakvik, Salvanes, and Vaage (2005), Humlum (2011)), but nevertheless they estimate very restrictive models. In particular, all these papers estimate regressions of child outcomes on the income of parents at different ages. Since the levels of income in different periods enter in a linear and additive way in these models, they are assumed to be "substitutes" in the production of human capital.

Relative to the papers using US and German survey data, our paper relies on much better data (larger samples and richer income histories), which allows us to estimate much more flexible models with considerable precision. This is also true even when we contrast our analysis to the ones using register data for Norway and Denmark. The flexible models we estimate enable us to construct a richer picture of the role of the timing of income than the one presented in previous work. This is important because the complementarity (or other interactions) of investments in human capital across periods (Cunha, Heckman, and Schennach (2010)) may translate into complementarity (or other interactions) of income shocks across periods.[4]

In another related paper, Caucutt and Lochner (2012) present an overlapping generations model of parental investments in children. In addition, Cunha (2013) also presents such a model

[4]Note that in a complementary paper, Hilger (2016) exploits parental layoffs to estimate the effect of parental job loss when children are aged 12-29 on children's long-term outcomes including college enrollment, earnings to age 25, but does not distinguish between the effects of different timings of this shock.

with overlapping generations and multiple periods of parental investments, explicitly accounting for uncertainty in income. The model we consider (and discuss in detail in an appendix) is much simpler than the models in any of these two papers. It is neither a general equilibrium model, nor an overlapping generations model, but it serves much more modest purposes including to 1) examine whether it is possible to explain the main patterns in the data using a parsimonious economic model of parental investments; 2) conduct simple simulations of the how the insurance provided by the tax and benefit system can influence the effect of the timing of family income on child outcomes.

The structure of the paper is as follows. In section 2 we discuss the specifics of the Norwegian welfare state, in section 3 we describe the data and in section 4 we present our empirical methods. Section 5 discusses our empirical results. Section 6 provides some interpretation of our results, including a brief discussion of problems caused by the potential endogeneity of the timing of income, the role of insurance possibilities, and a simulation of simple dynamic models of parental investments in children. Finally, section 7 concludes.

## 2 The Norwegian Welfare State

The Norwegian welfare state – together with the welfare states in Sweden and Denmark - has social insurance and economic equality as core elements. The Scandinavian model has been slowly developed over time, where the most important developments took place between World War II and the 1990s. For our purpose of understanding the effect of income shocks on children and teenagers, we will provide an overview of important aspects of the social security and transfer system and family policy.

Our sample period starts with births in 1970 and measures the latest child outcomes in 2014. Across this period, the Norwegian welfare state has consistently been characterised by highly progressive taxation, an extensive transfer system to redistribute income, and a strong social insurance system. Social security is characterized by universal benefits which cover the whole population rather than means testing, with many benefits and transfers related to current or previous labor force participation (Barth, Moene, and Willumsen (2014); Lindbeck (1997)).

The social security system, disability insurance and sick leave programmes are all based on income replacement whereby a worker is entitled to a benefit of a fixed percentage of the previous year's pay, or of the average over the last three years. (See Markussen and Røed (2014); Autor, Duggan, Greenberg, and Lyle (2016) for details on disability insurance and sick leave programmes).

A policy area which has seen reform in our sample period is family and active labour market policies, including maternity leave and childcare. The first maternity leave reform took place in 1977, providing full replacement for working mothers for four and half months (Carneiro, Løken, and Salvanes (2015)). This reform was extended successively in the 1980s and 1990s up until today with access to 49 weeks with 100 percent coverage (Dahl, Løken, Mogstad, and Salvanes (2016)).

Universal and strongly subsidised childcare was established in the late 1970s for children aged one to five years, but really expanded in the 1980s and 1990s (Havnes and Mogstad (2011); Black, Devereux, Løken, and Salvanes (2014)). Therefore, this reform is unlikely to have affected children from our sample, who were born in the period 1971-1980.

In addition to the family policies, an important area of change in our data period comes from changes in the progressivity of the taxation of income. The tax system has become less progressive over time through changes in the level of deductions, and especially of surtaxes (Thoresen, 2004). The major changes took place in 1992, 1995, 1999, 2001, 2004 and 2006 (Fjærli and Aaberge (1999); Blundell, Graber, and Mogstad (2015)). One example is the substantial tax reform of 1992, when personal income taxation was changed from a system where the nominal tax rate on income varied from 26.5 percent up to more than 40 percent, to a dual system where there was still progressive taxation of labor income, but capital income was taxed at a flat nominal rate of 28 percent. Moreover, the tax base was increased such that there was less room for deductions such as those on mortgage interest payments. There were several other changes as well as subsequent tax reforms, which together lead to a reduction in progressivity of the system as a whole. As documented in Blundell, Graber, and Mogstad (2015) the average tax rates were also strongly reduced across the period with the consequence of a reduction in the progressivity. Many of these changes were related to similar changes in all OECD countries in the same period (OECD (1998)).

# 3 Data

Our data source is the Norwegian Registry maintained by Statistics Norway, for the years ranging from 1971 up to 2014. It is a linked administrative dataset that covers the population of Norway, and it is a collection of different administrative registers providing information about (among other things) month and year of birth, educational attainment, labour market status, earnings, a set of demographic variables (age, gender), as well as information on families including parent's marital status. We are able to link individuals to their parents, and it is possible to gather labour market information for both fathers and mothers.

For the bulk of our analysis we select all births in the period 1971-1980. In particular, for our main analysis we construct annual household net earnings data for each year, starting from the three years preceding the child's birth through to their 20th birthday (we also construct earnings for the entire life cycle of parents which are used in our sensitivity analyses). The net income measure is calculated net of a progressive income tax, cash transfers, child benefits (for children up to age 18), unemployment and sickness benefits, general deductions for expenditures going to work and a regional compensation for living in the Northern most region. For some sensitivity analysis we analyse the effect of the timing of gross income, which is measured as pensionable income taking out sickness and unemployment benefits. This gives us information on income data and parental characteristics, mapped to children's outcomes, for 499,442 children.

We deflate income to 2000 and discount all incomes to the year of birth of the child, using a real interest rate of 4.26% (Aakvik, Salvanes, and Vaage (2005)). However, our results are robust to a large range of fixed discount rates between 0% and 15%, and to time-varying discount rates (matching the risk free interest rate evolution during the relevant years for our analysis). In order to construct a measure of income in each of the three periods we take the sum of discounted annual household incomes within each period (0-5, 6-11, 12-17). Permanent income is then defined as the sum of discounted income in the three periods.

Our data includes a large range of human capital outcomes including years of education for each individual, an indicator for dropping out of high school at the age of 16, and college enrollment. The consequences of the early dropout is that individuals do not receive a certificate for vocational or academic achievement which severely limits opportunities in the labor market, and prohibits access to further education. Unfortunately, it is not possible to measure whether a college degree was completed. Individuals are at least 26 years of age when we observe their educational achievement, and consequently, most of them can be expected to have exited school.

We also present results in the appendix for other outcomes including earnings at age 30, IQ and teenage pregnancy. Military service is compulsory in Norway for males, and between the ages of 18 and 20 males usually take an IQ test. This test is a composite of arithmetic, words,[5] and a figures tests[6], all of which are recognized as tests of IQ. For teen pregnancy, we construct an indicator for whether an individual (female) has a child, when she was between the ages of 16 and 20.

Finally, we use a detailed set of control variables, which are important to ensure that we are comparing differing income profiles across otherwise similar households. First, we construct household-specific slopes of income profiles by taking the difference between household income when the child is aged 18-20 and in the three years before birth.[7] Other controls include the child's year of birth and gender, as well as parental years of education and age at birth. Family size and marital status are endogenous so are not included in our benchmark specification but rather as a robustness check.

The descriptive statistics for the sample are reported in Table 1. There are 499,442 child-level observations, which include all individuals born in Norway between 1971 and 1980 for whom we were able to collect household income data. The permanent income of the household (in the period between the ages of 0 and 17 of the child) is about £312,200. There is substantial income dispersion: the standard deviation is £93,500. Income in each period (1, 2, and 3) falls with the age of the child because of discounting.

Mothers have on average 11.15 years of schooling, which is slightly lower than the average years of education of the fathers (11.47). Mothers are much younger than fathers at birth (26 vs 29 years

[5]The word tests are most similar to the Wechsler Adult Intelligent Scale (WAIS).

[6]The figures tests are similar to the Raven Progressive matrix.

[7]A robustness check conditions instead for the growth rather than the level difference of pre- and post- childhood income.

of age). The average years of education of the children in our sample is 12.74. 21% of children drop out from high school, but 40% attend college. The average annual earnings of these children at age 30 is £21,376. Regarding additional child outcomes, IQ is only available for males, and takes values on a 9 point scale, with a sample average of 5.26, and a standard deviation of 1.79. Finally, teen pregnancies occur for 8% of the females in our sample.[8]

# 4 Methods

Let $Y_i$ be an outcome for child $i$ (education, high school completion, college attendance, earnings, IQ, teenage pregnancy), measured in late adolescence or young adulthood. We are interested in $Y_i$ as a function of the history of (discounted) household income $I_{it}$ in each period $t$ ($t = 1, 2, 3$), and permanent income of the parents, $PI_i$. Since $PI_i = I_{i1} + I_{i2} + I_{i3}$ we drop one of the periods from the model, say $I_{i1}$. Therefore, we write:

$$Y_i = m(PI_i, I_{i2}, I_{i3}) + Z_i\delta + \varepsilon_i \tag{2}$$

$Z_i$ denotes a large set of controls which aim to capture that parents facing different income profiles may also be different in other dimensions. We condition on paternal and maternal education interacted with paternal or maternal age at birth (by including dummies for years of education and age at birth for each parent interacted with each other). This controls for different age-education profiles across households. Moreover, we construct a measure of household income growth between the ages of 0 and 17 of the child, based on income 1 to 3 years before birth (pre-birth income), and income 1 to 3 years after age 17 (post-17 income). This means that we explore fluctuations in income around deterministic age-income profiles which are allowed to vary with education, after accounting for heterogeneous income growth (and, of course, keeping fixed permanent income). The remaining controls in the model are birth year and gender of the child.

$\varepsilon_i$ should be interpreted as the unobserved heterogeneity that is left after controlling for permanent income and controls. We assume that $\varepsilon_i$ has finite conditional variance, $E(\varepsilon_i^2|PI_i, I_{i2}, I_{i3}, Z_i) \leq C < \infty$, and that $E(\varepsilon_i|PI_i, I_{i2}, I_{i3}, Z_i) = 0$. We are interested not in the impact of $PI$ itself on $Y$, but on the impact of the timing of income ($I_2$ and $I_3$) on $Y$, after conditioning on $PI$ and $Z$. In other words, we want to compare outcomes of children whose parents have the same level of permanent income between the ages of 0 and 17 and the same age-education profiles, household income growth, and child's birth year and gender, but differ in the level of income they experience in each period.

The assumption that $E(\varepsilon_i|PI_i, I_{i2}, I_{i3}, Z_i) = 0$ may be controversial. We would like to interpret

[8] The income process is studied in detail in Carneiro, Salvanes, and Tominey (2015). They find that, as in many other countries, the income process for Norwegian fathers can be fairly well approximated by the sum of a random walk (permanent shock) and a low order MA process (temporary shock).

$I_{i2}$ and $I_{i3}$ as income shocks orthogonal to other (unobservable) determinants of outcomes $Y_i$, conditional on $PI_i$ and $Z_i$. It is likely that $PI$ and $Z_i$ absorb much of the relevant unobserved heterogeneity across parents (correlated with the overall level of their income). We show in section 6.1 and Appendix section A3 that when we explore plausibly exogenous changes in the timing of income across households, resulting from changes in the tax and benefit system, our results are unchanged.[9]

We allow $m(PI_i, I_{i2}, I_{i3})$ to be a non-parametric function of its arguments. It tells us what counterfactual (adult) outcome $Y_i$ would be if a household faced a particular counterfactual income profile, $m(PI_i, I_{i2}, I_{i3})$. It is important to be able to estimate a flexible function, and the reason is the following. Parents are faced with income shocks in each period and in response, decide how much to invest in children (and how much to consume and save). There is a technology that links the adult human capital of an individual to the whole history of parental investments in childhood and adolescence. The link between income shocks and child outcomes, described by equation (2), depends on many factors, including preferences, technology, information, and the structure of credit markets (insurance possibilities). Therefore, there are likely to be complex nonlinearities and interactions between the different income measures included in the model.

We are particularly interested in $m_2(PI_i, I_{i2}, I_{i3}) = \frac{\partial m(PI_i, I_{i2}, I_{i3})}{\partial I_{i2}}$ and $m_3(PI_i, I_{i2}, I_{i3}) = \frac{\partial m(PI_i, I_{i2}, I_{i3})}{\partial I_{i3}}$. $m_2(PI_i, I_{i2}, I_{i3})$ tells us how the counterfactual outcome $Y_i$ for a child in a particular family would change (starting from some benchmark) if we reduced counterfactual income in period 1 and raised it in period 2, keeping $PI_i$ and $I_{i3}$ fixed (and $PI_i = I_{i1} + I_{i2} + I_{i3}$). In our empirical section we present a series of graphs relating $Y_i$ and $I_{i2}$ (for different outcomes $Y$). The graphs will vary depending on the values of $PI_i$ and $I_{i3}$ on which we evaluate this function. All other controls $Z_i$ are held at their mean values. An analogous interpretation and graphical representations of results can be given to $m_3(PI_i, I_{i2}, I_{i3})$.

Details of the non-parametric estimation methods are in Appendix A1.

[9] We also present results where we control for marital break up and number of children. In section 6.1 we repeat several of these arguments and provide a more detailed discussion of potential violations of the assumption that $E(\varepsilon_i | PI_i, I_{i2}, I_{i3}, Z_i) = 0$. In addition, one implication of our assumption is that pre-birth investments should be uncorrelated with the timing of income, but may still affect child outcomes. We test this in Appendix A3 by examining the relationship between having a low birth weight baby and the subsequent timing of parental income. We show that, although low birth weight is strongly correlated with $PI_i$, it is uncorrelated with $I_{i2}$ and $I_{i3}$. Similarly we show that income shocks occurring after the period in which child outcomes are measured do not drive those outcomes.

# 5 Results

## 5.1 Semi Parametric Estimates

In this section we present semi-parametric estimates of $m(PI, I_2, I_3)$, following the method described in Appendix A1.[10] The empirical results in this paper are presented through a series of graphs. We first focus on years of education as the outcome of interest, and then discuss results for earnings at age 30, IQ and teenage pregnancy. The outcomes high school dropout and college attendance are mentioned in the text but results are only presented in an Appendix. In order to implement the estimator we first need to create a grid of evaluation points for $m(PI, I_2, I_3)$. We take 19 points for each income variable $(PI, I_2, I_3)$, corresponding to the ventiles of each variables' distribution. This gives us a tri-dimensional grid with 6,859 points ($= 19 * 19 * 19$).

It is standard practice to trim the data so to avoid spurious results driven by small cells. Therefore, we drop 2% of the sample, corresponding to the cells with the smallest number of observations. In our main results we use a uniform kernel and choose the bandwidths using the formula in equation (1) of Appendix A1, setting $C = 6$. Below we show that our results are robust to the choice of different bandwidths.

One way to present our estimates of $m(PI, I_2, I_3)$ is through a series of two dimensional graphs, where in the y-axis we represent the outcome of interest, and in the x-axis we represent one of the income variables in $m(PI, I_2, I_3)$. The advantage of this presentation is that the graphs are straightforward to read. The downside of this type of presentation is that it allows variation of only one income period at a time, which means that we need to fix the remaining two income variables at some pre-determined values (we fix the remaining control variables at their mean values). Therefore, we create multiple figures for each outcome, corresponding to different fixed values for the left-out income variables in each graph.

For each outcome, we present three graphs. In the first graph, we fix $PI$ and $I_3$ at three different values each (the third, fifth, and seventh deciles of the distribution of each variable), and vary only $I_2$, for a total of nine possible combinations. These are presented in one figure, which plots $m(PI, I_2, I_3)$ against $I_2$ (for given values of $PI$ and $I_3$).

Since $PI = I_1 + I_2 + I_3$, when $PI$ and $I_3$ are fixed at particular values, we cannot vary $I_1$ and $I_2$ independently. Therefore, by moving towards the right along the x-axis in each graph, we are able to see how the outcome varies as $I_1$ decreases and, simultaneously, $I_2$ increases (i.e., as income is reduced in period 1 and increased period 2). In other words, our estimates concern the impact of changes in counterfactual family income profiles across different pairs of periods, holding constant permanent income and one per-period income ($I_3$ in this particular case) . The support of $I_2$ over which we can evaluate $m(PI, I_2, I_3)$ is not the same across all graphs within a figure because, for different combinations of $PI$ and $I_3$, there are values of $I_2$ for which there are no observations in

[10] Appendix A2 estimates the effect of the timing of income in a parametric setting, with corresponding Figures A1-A2.

the sample.[11] The second figure keeps $PI$ and $I_2$ fixed, and lets $I_3$ vary (so that income is reduced in period 1 and increased in period 3). The third figure keeps $PI$ and $I_1$ fixed and varies $I_3$ (a reduction in period 2 income and simultaneous increase in period 3 income).

In order to understand whether the patterns we uncover are statistically meaningful, below each figure we display two other parameters, $\alpha_1$ and $\alpha_2$ (and respective standard errors), which are defined as follows. For each panel, let $H$ be the highest grid point for the income variable being used in that panel, let $L$ be the lowest grid point, and $M$ be the median grid point (corresponding to the 50th percentile of the distribution of that income variable in that panel). Take the case where we fix $PI = \overline{PI}$ and $I_3 = \overline{I_3}$, and we let $I_2$ vary. Then define:

$$\begin{aligned} \alpha_1 &= m(\overline{PI}, M, \overline{I_3}) - m(\overline{PI}, L, \overline{I_3}) \\ \alpha_2 &= m(\overline{PI}, H, \overline{I_3}) - m(\overline{PI}, M, \overline{I_3}). \end{aligned} \quad (3)$$

$\alpha_1$ is the difference between the values the outcome takes in the median and lower extreme of the support of $I_2$, while $\alpha_2$ is the difference between the values the outcome takes in the median and upper extreme of the support of $I_2$. If $m(\overline{PI}, I_2, \overline{I_3})$ did not vary with $I_2$ (in which case the timing of income would be irrelevant, at least when we compare first and second period incomes) we would expect $\alpha_1 = \alpha_2 = 0$, so these parameters help us quantify the importance of the timing of income. Although there are several other ways to assess whether the curves we estimate are statistically different from what we would obtain if the timing of income was irrelevant (a flat line), the one we present here is particularly simple to implement and understand.

### 5.1.1 Years of schooling

We start by reviewing the evidence of the effect of the timing of parental income on years of schooling of the child. In panel a) of Figure 1 we represent the relationship between years of schooling of the child ($Y_i$) and $I_2$ keeping fixed $PI$ and $I_3$ (which means that $I_1$ decreases when $I_2$ increases, so we are measuring the impact on years of schooling of a counterfactual experiment which raises $I_2$ whilst simultaneously decreasing $I_1$). Panel b) represents the relationship between $Y_i$ and $I_3$ keeping fixed $PI$ and $I_2$ (impact of an increase in $I_3$ and simultaneous decrease in $I_1$), and Panel c) represents the relationship between $Y_i$ and $I_3$ keeping fixed $PI$ and $I_1$ (impact of an increase (decrease) in $I_3$ ($I_2$)).

Each panel represents 9 functions, which correspond to 9 different points of evaluation for 2 of the income variables omitted from the figure. For example, in Panel a) we keep $PI$ and $I_3$ fixed (while $I_2$, and therefore $I_1$ are allowed to change) at the 3rd (low), 5th (medium) or 7th (high) deciles of the distributions of each of these variables (since there are 2 omitted variables and 3 evaluation points for each of them, we get a total of $3^2 = 9$ curves). This allows us to see how the

[11] Note that we have data on the universe.

relationship between $Y_i$ and $I_2$ varies when we change the values of $PI$ and $I_3$.

It is useful to discuss one curve in one panel in detail. Then one can easily understand how to read all curves across all panels. Take the solid line in Panel a), corresponding to $PI$ low and $I_3$ low. At the further left point of this line, corresponding approximately to $I_2 = \pounds 60000$, the average child has about 12.5 years of schooling. Since $PI^{low} = \pounds 250930$ and $I_3^{low} = \pounds 81000$, if $I_2 = \pounds 60000$ then $I_1 = \pounds 250930 - \pounds 60000 - \pounds 81000 = \pounds 109930$. As $I_2$ increases from £60000 to £120000 or so (and $I_1$ decreases from £109930 to £49930), $Y_i$ falls by about half a year of schooling from about 12.5 to below 12.

Within each panel of Figure 1, even though different curves are evaluated at dramatically different values of the omitted income variables and clearly have different locations, they display strikingly similar shapes. These similarities can be better seen in Figure 2, where we adjust all the curves in each panel to have a similar location. More precisely, in Figures 2a-c) the graphs have been adjusted along the y-axis to have a common mean value of the child outcome, and in a small number of curves there was also an adjustment along the x-axis so that the peak of each curve occurs at a similar value of x.[12] The adjustments are additive (adding constants horizontally and vertically to the curves in Figure 1) so they do not influence the curvature of the functions shown here. In the rest of the paper we will present graphs with similar adjustments to those in panels a)-c) of Figure 2 (the appendix shows the full set of figures before adjusting the axes).

To test whether the patterns in the figures are statistically meaningful, we compute $\alpha_1$ and $\alpha_2$ (and report standard errors for the test that these parameters are equal to zero). There is one $\alpha_1$ and one $\alpha_2$ for each curve in each panel of Figures 2, so for each panel we report the average values of $\alpha_1$ and $\alpha_2$ across the nine different curves, and the corresponding standard errors (the individual values of these parameters are shown in Figures A3ai-jiii) in the appendix).

All of the graphs in Figure 2a) (and Appendix Figures A3ai-ciii) are quite similar. The dominant pattern is a downward sloping curve, whereby an increase in $I_2$ and a decrease in $I_1$ lowers years of schooling. The slope of this curve is statistically different from zero (especially when we look at $\alpha_2$). This suggests that income occurring at ages 0-5 has a relatively larger impact on years of schooling than income occurring during ages 6-11.

For very low levels of I2 (very high levels of $I_1$), an increase in $I_2$ initially raises years of schooling in some of these figures. Close inspection of $\alpha_1$ in Figure 2a) (and Appendix A3ai-ciii) however tells us that this initial increase is not generally statistically different from zero.[13]

[12] Specifically, for Figure 2b) one curve was shifted by a value of 1 to the right (PI and I2 low) and two curves shifted by the value of 1 to the left (PI med; I2 high and PI high; I2 high).

[13] As a benchmark, it is useful to compare the downward slopes in the different panels of Figure 2a), with the slope of the relationship between years of schooling and permanent income, shown in Appendix Figure A1. The slopes of the curves in Figures 2a) are of a very similar magnitude of those shown in Figure A1. For example, a £50,000 increase in permanent earnings leads to about an extra 0.5 years of education, while an increase in middle childhood's income from £80,000 to £130,000 leads to a change in schooling of the same magnitude, of about 0.50 years of education. It is remarkable that, even conditional on permanent income, the impact of the timing of income is as large as the impact of permanent income.

The slope of the curves in Figures 2a) become steeper as the level of I1 falls, as evidenced from a comparison of $\alpha_1$ and $\alpha_2$. $\alpha_1$ takes the value of -0.017 suggesting a reduction of $I_1$ and simultaneous increase of $I_2$, from the bottom to the middle of the distribution of $I_2$, lowers schooling by 0.017 years. The corresponding figure for $\alpha_2$ suggests that a shift from $I_1$ to $I_2$, starting from the midpoint of $I_2$ to the highest level lowers schooling by 0.585 years. This says that the larger impact of income at 0-5 compared to income at 6-11 is especially large when $I_1$ is low (rather than when it is high).

Consider now the productivity $I_1$ relative to $I_3$. In Figures 2b) (and Appendix Figures A3di-fiii) $I_2$ and $PI$ are fixed within each graph at the third, fifth and seventh decile, and we show how the estimate of $m(\overline{PI}, \overline{I_2}, I_3)$ increases with $I_3$ (corresponding to decreasing levels of $I_1$). There is an inverse-U shape pattern for nearly all the curves displayed in the figure. Years of schooling initially increase as income is moved from (initially high) $I_1$ to (initially low) $I_3$, reach a maximum at around £100,000 of (discounted) $I_3$, and fall thereafter. $\alpha_1$ is positive and statistically different from zero, and $\alpha_2$ is negative and statistically different from zero.

What this means is that, across different values of $PI$ and $I_2$, the years of schooling of the child are maximized when there is some balance between early and late childhood income. If income is too concentrated in either early or late childhood, then one can improve education outcomes of children by "shifting" income towards the other period. Looking at the estimate for $\alpha_1$, moving income from a low initial level of $I_3$ (and a high level of $I_1$) to the median level of $I_3$ raises years of schooling by 0.324 years. Looking at $\alpha_2$, moving income from the median level in period 3 towards very high levels of $I_3$ (and consequently very low levels of $I_1$) reduces years of schooling by 0.162 years.

Figure 2c) (and Appendix Figures 3gi-jiii) examines the trade-offs between income in middle and late childhood (keeping fixed income in early childhood and permanent income). The figures display an increasing shape, indicating that years of schooling are maximised when income is reduced in middle childhood and increased in late childhood. Across figures, $\alpha_1$ and $\alpha_2$ are both positive and statistically different to zero. The magnitudes of these impacts are substantial and similar, if slightly smaller than comparable changes in permanent income. The slopes of the estimated curves are particularly steep when $I_3$ is increased from initially low levels, where $\alpha_1$ and $\alpha_2$ take the value of 0.313 and 0.201 respectively.

It is striking how similar the displayed patterns are across the different curves, within each panel of Figure 1 and 2. We have essentially decreasing functions when looking at $m(\overline{PI}, I_2, \overline{I_3})$, inverse U-shapes when looking at $m(\overline{PI}, \overline{I_2}, I_3)$, and increasing functions when displaying $m(\overline{PI}, \overline{I_1}, I_3)$. We do not impose this in our estimation procedure, and were surprised that we did not have more variety in the estimated patterns, depending on, for example, how wealthy the family was (in terms of permanent income). Our initial goal was to have a framework flexible enough to uncover diversity in the behavior of these functions, but what we uncovered was instead stability. The stability in the shapes of these functions is also observed when we study other outcomes, as we show below.

These results we show have three important implications. First, the timing of income shocks is relevant for human capital formation. If the timing of income was irrelevant then all these graphs would be horizontal lines, with only permanent income being relevant for human capital outcomes. It is likely that the timing of income shocks affects the timing of investments in human capital (which in turn affects human capital formation). This will happen if parents have imperfect insurance possibilities against income shocks.

Second, if the timing of income matters because parental investments react to income shocks, the shape of the curves in Figures 2b) is consistent with an underlying technology of skill formation that exhibits complementarity in investments across periods. The reason is that, under complementarity, it is desirable to maintain a balanced flow of investments over the life of the child. However, at this point, we cannot distinguish this model from any other model where smoothing investments is desirable (see the discussion in Appendix A6).

Third, our results are also consistent with a model where investments in the second period of childhood (ages 6-11) have relatively lower productivity than investments in early childhood and in late childhood. Given that all our outcomes are measured in late adolescence and late adulthood, it is perhaps not surprising that later investments have "higher productivity" than earlier investments, just because they had less time to depreciate. If anything, what is surprising is that early investments can be so much more important than investments in middle childhood, and as important as investments in late childhood. This may happen because they depreciate less than later investments, and/or because they are complementary with investments in late childhood.

It is important to understand the upward sloping section in each curve (see also Appendix Section A5). One might think that, for a given level of permanent income, it should not be worse to have all income available in the first period (or in an earlier period), than to receive it in payments spread out over different periods of childhood. This is because, if permanent income becomes fully available at time zero, then one can allocate it freely across periods just by saving the appropriate amount, regardless of whether or not one can borrow. If this is true then children in households getting all their income in early childhood (and nothing in later period) should not do worse than children in households with a more spread apart flow of income (keeping permanent income fixed).

However, the reasoning just described ignores the fact that incomes are uncertain, and parents cannot fully anticipate future shocks. They cannot be certain of whether their income will rise or decline in future periods. In addition, their objective function is not just maximizing the net present value of child outcomes, but they trade off their own consumption for investments in children (unless capital markets are perfect). As a result, in the presence of income uncertainty, when faced with income shocks parents might oversave for precautionary motives and invest in children less than they could, unless they have perfect insurance. An income shock occurring late in the life cycle can induce additional investments in this period of life that might not have occurred if the shock took place earlier in life. These investments could in turn lead to an improvement in child outcomes.

### 5.1.2 Gender Differences

There is a large literature documenting gender differences in parental investments in child human capital in developing countries (e.g., Pitt and Rosenzweig (1990); Rosenzweig and Schultz (1982); Parish and Willis (1994); Quisumbing (1994)). In developed countries, parental investments and the returns to these investments have also been found to be different between boys and girls. Baker and Milligan (2016) document higher levels of parental investments which promote learning in girls compared to boys, which lead to a higher score for reading and maths achievement of girls compared to boys. Bertrand and Pan (2013) find that the response of children's noncognitive skills to parental investments are much more sensitive for boys than girls, suggesting that the production function which translates investments to child outcomes may vary across child gender.

In Figures 3a-c) (and Appendix Figure A4) the semiparametric estimates of the effect of the timing of income on years of schooling are shown for girls, while estimates for boys are in Figures 4a-c) (and Appendix Figure A5). The figures for girls are similar to those reported above, especially in panels a) and c). There is a less pronounced inverse-U shaped pattern for panel b). In contrast, for boys there is a more pronounced inverse-U shaped pattern panel b) (concerning the productivity of $I_3$ relative to $I_1$) than when we looked at the whole sample, and several curves in panel a) (concerning the productivity of $I_2$ relative to $I_1$) also display a strong inverse-U shape (confirmed by the estimates of $\alpha_1$ and $\alpha_2$ in that same panel). Panel c) of Figure 4 is similar to that of Figures 2 and 3.

As explained above, one possible explanation for inverse-U shaped patterns in the relative productivity of income in one age vs another is consistent with a model where investments in children react to income shocks, and there is dynamic complementarity between investments in different time periods. Our estimates suggest that, for boys, this complementarity is more important than for girls.

### 5.1.3 Other outcomes: earnings at age 30, IQ, teenage pregnancy

In addition to years of schooling, we considered high school dropout rates and college attendance rates separately, since they correspond to education decisions at the lower tail and at the upper tail of the education distribution. Perhaps not surprisingly, these results are broadly similar to the ones we showed for years of education, and we display these results in the Append (Figures A7-A10). More interestingly, we examine what our results look like when we consider three other outcomes for the child: log earnings measured at age 30, IQ, and teenage pregnancy. We turn to these next.

Most of the literature on intergenerational mobility links the earnings of parents to the earnings to children. Figures 5a-c) (and Figure A6ai-jiii in the Appendix) show semi-parametric estimates analogous to those discussed above (e.g., Figure 2), but using as outcome the log annual earnings of children at age 30. Figures 5a) (Figure A6ai-ciii in the Appendix) shows a downward sloping curve ($I_2$ vs $I_1$), Figures 5b) (Figure A6di-fiii in the Appendix) show an inverse u-shaped curve ($I_3$ vs $I_1$), and Figures 5c) (Figure A6gi-jiii in the Appendix) show an upward sloping curve ($I_3$ vs $I_2$). These are remarkably similar to the schooling results, suggesting that whatever component of human capital is driving the schooling results, may also be driving what we observe for earnings.

Figures 6 and 7 (Figures A11-A12 in the Appendix) show our estimates when we use IQ and teenage pregnancy as the outcomes of interest. It is important to look at outcomes such as these because they could potentially reflect different dimensions of human capital than those manifested through schooling and earnings. IQ is measured through the armed forces tests at age 18, only for males, and is traditionally seen as a good measure of cognitive ability. Teenage pregnancy, measured only for girls, may reflect instead behavior problems, excessive risk taking, or low ability to plan for the future. It is again remarkable that, with one exception, the patterns we uncover for the impact of the timing of income on these measures of child human capital are very similar compared to years of schooling and earnings at age 30.

Figure 6 concerns IQ. Panels a) ($I_2$ vs $I_1$) and c) ($I_3$ vs $I_2$) are similar to those seen before for schooling and earnings, showing that $I_1$ and $I_3$ are both relatively more productive than $I_2$. Panel b) ($I_3$ vs $I_1$) shows mainly increasing curves (whereas we saw inverse-U shaped curves for schooling and earnings, suggesting that $I_3$ is relatively more important than $I_1$. This is the only instance where there is a qualitative difference in the patterns we uncover for schooling, earnings, IQ and teenage pregnancy.

Figure 7 concerns teenage pregnancy, and shows qualitatively the same patterns we saw for schooling and earnings, although the signs of the slopes are reversed since we see teenage pregnancy as a reflection of low human capital. The U shaped pattern between $I_1$ and $I_3$ is quantitatively less pronounced than (the inverse-U shaped pattern) for other outcomes, but it is still statistically significant. In sum, we see broadly similar results when we study the relationship between the timing of family income and very different child outcomes.

# 6 Interpretation of our results

## 6.1 Exogeneity of the Timing of Income Shocks

In our main empirical specification we cannot guarantee that the error term is orthogonal to the timing of income, even after controlling for permanent income. There are challenges to interpreting our estimates of the impact of the timing of income on the human capital of children as causal. We discuss these briefly in this section, and present a much more detailed treatment of this issue (and of all the results summarized below) in Appendix A3.

First, the precipitant of the income shock that we estimate and not the income shock itself may be the driver for the estimated effects on child outcomes. One particular example of this issue is related to labour supply. The decision of a parent to exit the labour market and spend more time with her children could lower household income and directly raise child outcomes. We show however that the variation in the timing of income does not come from participation decisions or hours worked. Only 0.42% and 1.46% of household and fathers' net income, respectively, is explained by labour market participation, and similarly, only 3.4% and 3.2% of the variation in net household and fathers' income, respectively, is explained by variation in hours worked (see section A3.1 in Appendix 3).

Second our results are robust to using exogenous variation in the timing of income generated by changes in the mapping between gross and net income, caused by reforms to the tax and benefit system occurred between 1971 and 1998 (e.g. changes to the thresholds of tax brackets). Changes in the tax and benefit system means that two individuals with the same gross income profile may have very different net income profiles depending on their cohort. Such variation in the timing of net income, conditional on gross income, is plausibly unrelated to unobservable traits correlated with the timing of parental income and child outcomes, where identification comes from the cohort variation in the exposure to these changes (see section A3.1 in Appendix 3).

Third, one could worry that our finding that middle income is relatively less productive than income in later adolescence is driven by the omission from our model of permanent income after the age of 17, which is highly correlated with income in adolescence. For example the ability of parents to borrow against future income (say between the child age 18-30) is likely to be higher at age 12-17 compared to age 6-11. However, there is hardly any change to our results when we carefully expand the model to include the whole lifecycle of families, either by taking family income for the periods the child is between ages 0 and 30, or for the periods when parents are between ages 25 and 55 (see section A3.2 in Appendix 3).

Fourth, the timing of income could be a choice variable, potentially correlated with human capital investments. One obvious case is maternity leave: parents choose to take time off at the beginning of the child's life, presumably with the goal of improving the outcomes of the child (e.g., Carneiro, Løken, and Salvanes (2015)). However, excluding income in the early years when the labour supply of mothers (and to a lesser extent fathers) are affected by child-rearing activities leads to the same conclusions as our benchmark. We also conduct two placebo tests which further reinforce our belief that, in our very rich specification, we can consider the timing of income as exogenous. We show that the timing of income does not matter for low birth weight - an outcome which predates the income. In addition, we also show that shocks to parental income occurring when the child was 30 years of age do not affect her education outcomes (see section A3.3 in Appendix 3).

Finally, there may be unobserved traits of individuals which are potentially correlated with human capital investments and with income profiles. Those parents who have their income front-loaded in the early years of the life of the child could be different from those who have their income back-loaded. The question is, in which ways are they different?[14] Our results are robust to several checks, including using different measures for the income growth profile and controlling for income volatility (see sections A3.4 and A3.5 in Appendix 3).[15]

## 6.2 Insurance

The timing of income is important in a world where households are not able to fully insure themselves against unanticipated income shocks[16] As shown in Cunha and Heckman (2007) and Caucutt and Lochner (2012), if families face lending or borrowing restrictions, then two families with the same statistical permanent income but different time profiles of income may also have different time profiles of investments in children, resulting in children with different levels of human capital. This result, which is the basis for the empirical procedure implemented in this paper, is consistent with the standard dynamic model of parental investments in children with uncertainty.

Beyond self insurance possibly available through access to credit markets, there are other types of insurance mechanisms available to households that we can think of being particularly relevant to our analysis. The most obvious one is the tax and benefit system, which in Norway is quite progressive. This implies that parents are more likely to be insured against gross income shocks than net income shocks, and that the timing of parental gross income should matter less than the timing of parental net income. Blundell, Graber, and Mogstad (2015) find that the tax and benefit system provides substantial insurance against gross income shocks. Since all our previous estimates are obtained using net household incomes, we repeated our analysis using gross incomes, and found that the resulting curves are generally flatter than the curves in Figure 2, suggesting that the tax and benefit system does provide some insurance (see Figures A35 in the Appendix). The trade offs between $I_2$ and $I_1$, and between $I_3$ and $I_2$, are definitely flatter when we look at gross income profiles (Figure A35 in the Appendix) than net income profiles (Figure 2 in the paper). In fact the relationship between child's schooling and $I_2$ relatively to $I_1$ is even inverse-U shaped when we consider gross income, rather than declining when we use net income (perhaps because progressive income taxation means that a unit of gross income at very low level of income is worth more than a unit of gross income at a high level of income). This can be seen clearly by comparing $\alpha_1$ and $\alpha_2$ in Figures A35, which represent the difference between schooling when income moves from the bottom to the midpoint of the distribution, and from the midpoint to the top of the distribution, respectively. Both $\alpha_1$ and $\alpha_2$ are lower when estimating the effect of gross income. The trade off between $I_3$ and $I_1$ is qualitatively and quantitatively more similar regardless of using

[14]Parental education affects the slope of income profiles and the outcomes of children, but this is a variable we control for.

[15]Our results are robust to several specification checks, which vary both the set of controls, the measurement of controls, the discount rates to construct per period income and the choice of bandwidth in the nonparametric estimation.

[16]Appendix Section A4 provides full discussion of the analysis.

gross or net income, which would suggest that the tax and transfer system does not provide too much insurance to this trade-off, perhaps because it is difficult to insure between periods that are well apart from each other than periods that are close together.

It is also interesting to consider the potential role of spousal insurance by comparing results using father and household incomes. If spousal insurance was important then we would estimate larger effects of household income shocks relative to father income shocks, since the former cannot be insured through this mechanism. This is in fact what we find (see Figures A36 in the Appendix, which have the same shapes but less pronounced slopes than Figure 2 in the paper).

The importance of within family insurance is also visible when we compare the role of the timing of income across siblings within the same family. Parents may be able to partly insure investments in a specific child against negative income shocks affecting that child in a particularly sensitive period, by shifting investments away from other children in the household. Indeed, when we use only within sibling variation we find no evidence of a relationship between the timing of income fluctuations and child outcomes (see Figure A37 in the Appendix). This suggests that there is substantial within-family insurance against income shocks, since our estimates using cross family variation in the timing of income shows much more variation in child outcomes.

## 6.3 Simulating Large Changes to the Tax and Benefit System

We estimated a simple model of parental investments in children to simulate the extent to which changes in the progressivity of the tax and benefit system influence the way income fluctuations over the life cycle of the child affect human capital formation. With a model, one can ask which level of progressivity would be necessary to completely smooth out the effect of the timing of income shocks and how it compares with the current system and with large changes to the tax and benefit system in Norway over the last 40 years.

The structural model we consider is very simple, so we refrain from providing a detailed description here (it is described in detail in Appendix 5). There are 3 periods of childhood, parents optimally choose time and good investments in each period, as well as employment, in a world where there is imperfect insurance against income shocks, so that investments in children react to these shocks. We estimate the technology of skill formation, assumed to be a Constant Elasticity of Substitution (CES) production function, where the output is years of education of the child and the inputs are investments in each period. The estimated elasticity of substitution between investments in different time periods is such that investments in the first and third periods of childhood would be more productive than investments in the second period, supporting the idea of the presence of dynamic complementarities in investments (the model fits both patterns of human capital and labor market participation quite well).

Within this model we simulate how the timing of income affects human capital formation under the following alternative tax/benefit systems: i) a less progressive system than the Norwegian system, such as the United States; ii) a slightly more progressive system than Norway; and iii) a system with a very extreme degree of progressivity. We find that across a wide range of degrees of progressivity of the tax system, the timing of income shocks is always important, unless we make progressivity unrealistically large (as in case iii).

# 7 Conclusion

This paper examines the importance of the timing of income shocks for the human capital development of children. Using a very large dataset, consisting of the entire population of children born in Norway between 1971 and 1980, we estimate semi-parametric regressions of human capital outcomes of children (measured in their adult years), on the discounted household income for the years when the child was between 0 and 17 years of age (which we label permanent income), and on income in different periods of childhood.

Using a method which controls for permanent income across the entire period of childhood in essence compares one family with a counterfactual family whose permanent income is the same, but with a different configuration of the timing of income shocks. The paper asks whether child development responds to this differential timing of family income.

We show that, keeping permanent income and income in late childhood fixed, child outcomes such as years of schooling and earnings by age 30 improve if a family experiences relatively high (counterfactual) income in early childhood and low income in middle childhood, rather than a relatively low income in early childhood and high income in middle childhood. Stating it more simply, child outcomes improve as (counterfactual) family income rises in early childhood and falls in middle childhood. We also show that a child's schooling and earnings increase when her family income falls in the middle years and increases towards the later years of childhood. Compared to the early and late period of childhood, income in the middle period seems to have low productivity.

Finally, children in families with a balanced profile of income between early and late childhood (and keeping fixed permanent income) experience better adult outcomes than those in families with income mainly "front-loaded" in the early period, and those in families with income mainly "back-loaded" in the later period of childhood. In other words, we find that there is an inverse U-shaped relationship between income the early period relative to income in the late period (which means keeping fixed permanent income and income in middle childhood), and the child's education and labour market outcomes in the adult years.

These empirical patterns are consistent with a model with income uncertainty, partial insurance (Blundell, Pistaferri, and Preston (2008)), and complementarity between investments in children across periods (Cunha and Heckman (2007), Cunha, Heckman, and Schennach (2010)). Income uncertainty, together with partial insurance possibilities and the technology of skill accumulation, leads to a setting where investments in children react to parental income shocks. Complementarity between investments taking place in different periods means that human capital is maximized when there is a balanced flow of investments. The combination of these factors can result in a model where a balanced flow of income shocks may lead to higher human capital than an unbalanced history of shocks. This would be consistent with an inverse U-shaped relationship between the education of the child and the amount of income that is front-loaded (or back-loaded) in the initial (late) period of childhood. However, our results also suggest that income in the middle period of childhood may be relatively unproductive, compared to income in earlier and later periods.

# References


Aakvik, A., K. G. Salvanes, and K. Vaage (2005): "Educational Attainment and Family Background," *German Economic Review*, 6(3), 377–394.

Autor, D. H., M. Duggan, K. Greenberg, and D. S. Lyle (2016): "The impact of disability benefits on labor supply: Evidence from the VA's disability compensation program," *American Economic Journal: Applied Economics*, 8(3), 31–68.

Baker, M., and K. Milligan (2016): "Boy-Girl Differences in Parental Time Investments: Evidence from Three Countries," *Journal of Human Capital*, 10(4), 399–441.

Barth, E., K. O. Moene, and F. Willumsen (2014): "Reprint of "The Scandinavian model-An interpretation"," *Journal of Public Economics*, 127, 17–29.

Becker, G. S., and N. Tomes (1979): "Becker-Tomes_(JPolE_79).pdf," *Journal of Political Economy*, 87(6).

——— (1986): "Human capital and the rise and fall of families.," *Journal of labor economics*, 4(3 Pt. 2), 1–47.

Bertrand, M., and J. Pan (2013): "The Trouble with Boys: Social Influences and the Gender Gap in Disruptive Behavior.," *American Economic Journal: Applied Economics*, 5(1), 32–64.

Björklund, A., and K. G. Salvanes (2011): "Education and Family Background: Mechanisms and Policies," *in Hanushek, Eric, Stephen Machin and Ludger Woessmann, Handbook in Economics of Education, 201-247, North-Holland.*, 3(11).

Black, S. E., and P. J. Devereux (2011): "Recent Developments in Intergenerational Mobility," *in Handbook of Labor Economics, Volume 4B*, 4(11), 1487–1541.

Black, S. E., P. J. Devereux, K. V. Løken, and K. G. Salvanes (2014): "Care or Cash? The Effect of Child Care Subsidies on Student Performance.," *The Review of Economics and Statistics*, 96(5), 824–837.

Blundell, B. R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 0214491(1998), 1887–1921.

Blundell, R., M. Graber, and M. Mogstad (2015): "Labor income dynamics and the insurance from taxes, transfers, and the family," *Journal of Public Economics*, 127, 58–73.

Carneiro, P., and J. J. Heckman (2003): "Human Capital Policy," *in Heckman, J. J. and Krueger, A. Inequality in America: What Role for Human Capital Policies, MIT Press.*, (2000).

Carneiro, P., K. V. Løken, and K. G. Salvanes (2015): "A Flying Start? Maternity Leave Benefits and Long-Run Outcomes of Children," *Journal of Political Economy*, 123(2), 365–412.

Carneiro, P., K. Salvanes, and E. Tominey (2015): "Family Income Shocks and Adolescent Human Capital," *Unpublished Manuscript*.

Caucutt, E., and L. Lochner (2012): "Early and late human capital investments, borrowing constraints, and the family," *NBER Working Paper Series*.

Cunha, F. (2013): "Investments in children when markets are incomplete," *Unpublished Manuscript*.

Cunha, F., and J. Heckman (2007): "The Technology of Skill Formation," *American Economic Review*, 97(2), 31–47.

Cunha, F., J. Heckman, and S. Schennach (2010): "Estimating the Technology of Cognitive and Noncognitive Skill Formation.," *Econometrica*, 78(3), 883–931.

Dahl, B. G. B., K. V. Løken, M. Mogstad, and K. V. Salvanes (2016): "What is the Case for Paid Maternity Leave?," *Review of Economics and Statistics*, 98(2), 209–225.

Duncan, G. J., W. J. Yeung, J. Brooks-gunn, and J. R. Smith (1998): "How Much Does Childhood Poverty Affect the Life Chances of Children?," *American Sociologial Review*, 63(3), 406–423.

Fjærli, E., and R. Aaberge (1999): "Tax Reforms, Dividend Policy and Trends in Income Inequality: Empirical Evidence Based on Norwegian Data," .

Havnes, T., and M. Mogstad (2011): "No Child Left Behind: Subsidized Child Care and Children's Long Run Outcomes," *American Economic Journal: Economic Policy*, 3(May), 97–129.

Hilger, N. G. (2016): "Parental job loss and children's long-term outcomes: Evidence from 7 million fathers' layoffs," *American Economic Journal: Applied Economics*, 8(3), 247–283.

Humlum, M. K. (2011): "Timing of Family Income , Borrowing Constraints and Child Achievement," *Journal of Population Economics*, 24(3), 979–1004.

Jenkins, S. P., and C. Schluter (2002): "The Effect of Family Income During Childhood on Later-Life Attainment : Evidence from Germany," *IZA Discussion Paper Series*, (604).

Levy, D. M., and G. Duncan (2000): "Using sibling samples to assess the effect of childhood family income on completed schooling," Discussion paper, Northwestern University/University of Chicago Joint Center for Poverty Research.

Lindbeck (1997): "The Swedish Experiment," *Journal of Economic Literature*, 35(3), 1273–1319.

Markussen, S., and K. Røed (2014): "The Impacts of Vocational Rehabilitation The Impacts of Vocational Rehabilitation," *IZA Discussion Papers Series*, No. 7892(7892), 1–13.

OECD (1998): "OECD Economic Outlook," Discussion Paper June.

Parish, and Willis (1994): "Sons, Daughters, and Intergenerational Support in Taiwan," *American Journal of Sociology*, 99(4), 1010–1041.

Pitt, and B. M. R. Rosenzweig (1990): "American Economic Association Productivity , Health , and Inequality in the Intrahousehold Distribution of Food in Low- Income Countries," *American Economic Review*, 80(5), 1139–1156.

Quisumbing, A. (1994): "Improving Women's Agricultural Productivity as Farmers and Workers," *Education*, (July), 93.

Rosenzweig, B. M. R., and T. P. Schultz (1982): "American Economic Association Market Opportunities , Genetic Endowments , and Intrafamily Resource Distribution : Child Survival in Rural India," *American Economic Review*, 72(4), 803–815.

Solon, G. (1999): "Intergenerational Mobility in the Labor Market," *In Handbook of Labor Economics, Volume 3. O. Ashenfelter and D. Card, editors.*, 3, pp.1761–1800.

Figure 1: Semi-parametric Estimates. Dependent Variable is Years of Schooling.

a) Household Income 6-11

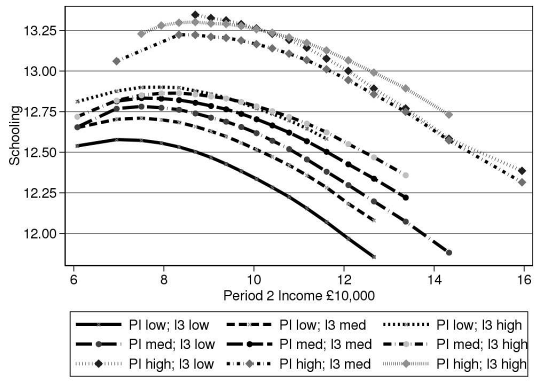


$\alpha_1 = -0.017\,(0.013)$ $\alpha_2 = -0.585\,(0.015)$

b) Household Income 12-17

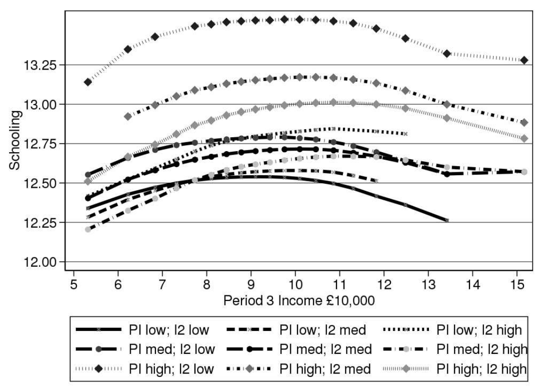


$\alpha_1 = 0.324\,(0.014)$ $\alpha_2 = -0.162\,(0.016)$

c) Household Income 12-17

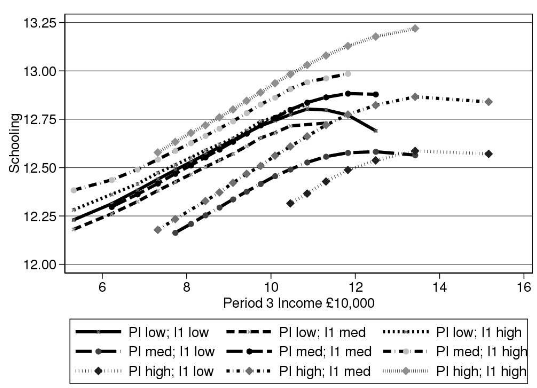


$\alpha_1 = 0.313\,(0.014)$ $\alpha_2 = 0.201\,(0.016)$

Note: Income in 2000 prices, £10,000s. Semi-parametric estimates control for dummies for paternal education interacted with age and maternal education interacted with age, household net income profile, gender and child year of birth and cluster at the family level. $\alpha_1$ and $\alpha_2$ and associated standard errors defined in equation 3 indicate the significance of the curvature, denoting difference in outcome between median and lower extreme, and between upper extreme and median respectively. Low, medium, high levels of permanent income (PI), I1, I2, I3 indicate income at the 3rd, 5th and 7th decile.

Figure 2: Semi-parametric Estimates. Dependent Variable is Years of Schooling: Adjusted Figures.

a) Household Income 6-11

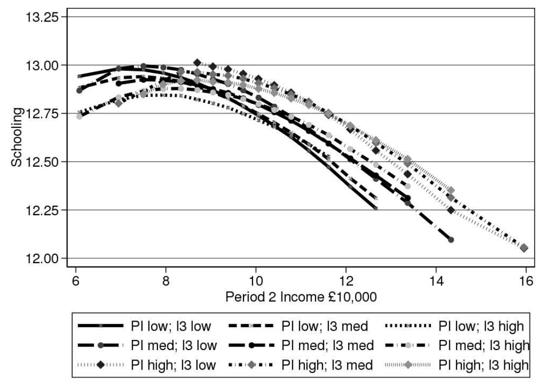


$\alpha_1 = -0.017\,(0.013)$ $\alpha_2 = -0.585\,(0.015)$

b) Household Income 12-17

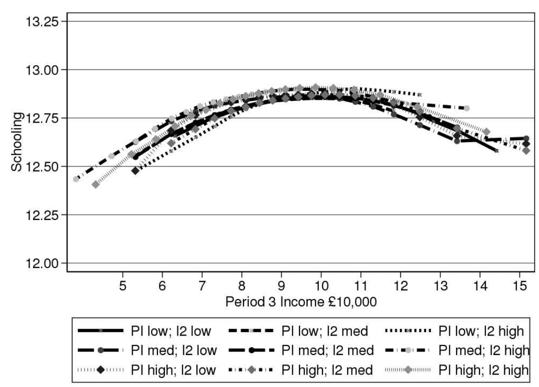


$\alpha_1 = 0.324\,(0.014)$ $\alpha_2 = -0.162\,(0.016)$

c) Household Income 12-17

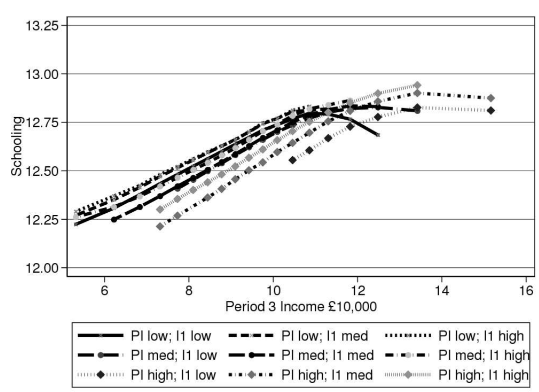


$\alpha_1 = 0.313\,(0.014)$ $\alpha_2 = 0.201\,(0.016)$

Note: Income in 2000 prices, £ 10,000s. Semiparametric regression controls and interpretation of $\alpha_1$ and $\alpha_2$ and associated standard errors is the same as in Figure 1. Low, medium, high levels of permanent income (PI), I1, I2, I3 indicate income at the 3rd, 5th and 7th decile. An adjustment of the scale of the x-axis and for 2b) the y-axis as described in Section 5.1.

Figure 3: Semi-parametric Estimates. Dependent Variable is Years of Schooling: Girls.

a) Household Income 6-11

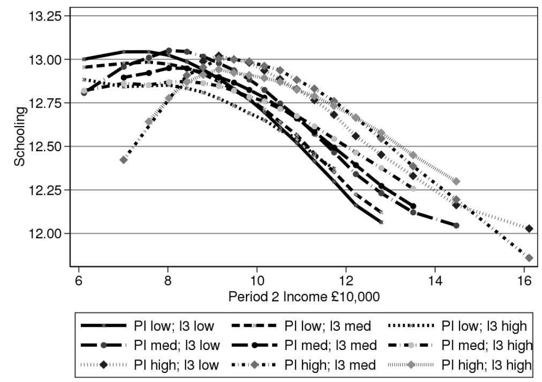


$\alpha_1 = 0.018\ (0.029)\ \alpha_2 = -0.708\ (0.033)$

b) Household Income 12-17

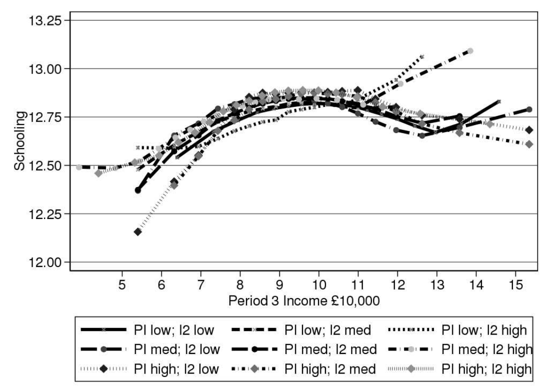


$\alpha_1 = 0.401\ (0.031)\ \alpha_2 = -0.014\ (0.038)$

c) Household Income 12-17

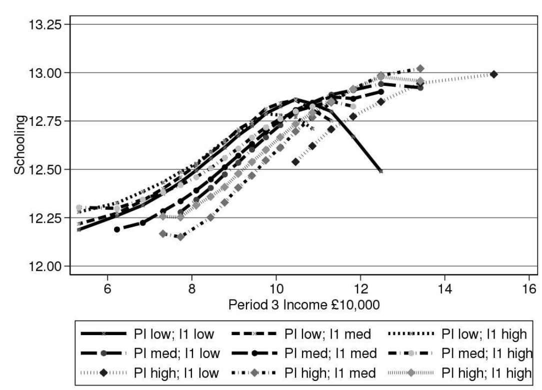


$\alpha_1 = 0.350\ (0.030)\ \alpha_2 = 0.222\ (0.033)$

Note: Income in 2000 prices, £ 10,000s. Semiparametric regression controls and interpretation of $\alpha_1$ and $\alpha_2$ and associated standard errors is the same as in Figure 1. Low, medium, high levels of permanent income (PI), I1, I2, I3 indicate income at the 3rd, 5th and 7th decile. An adjustment of the scale of the x-axis and for 3b) the y-axis as described in Section 5.1.

Figure 4: Semi-parametric Estimates. Dependent Variable is Years of Schooling: Boys.

a) Household Income 6-11

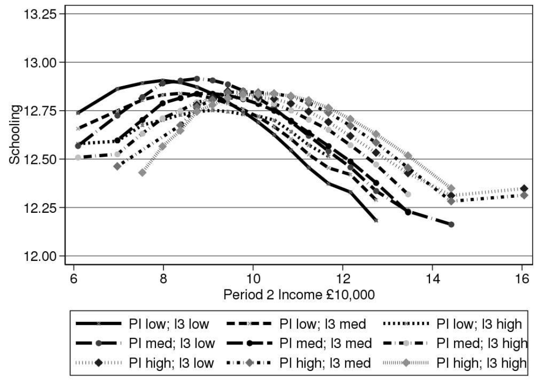


$\alpha_1 = 0.217\,(0.024)$ $\alpha_2 = -0.507\,(0.028)$

b) Household Income 12-17

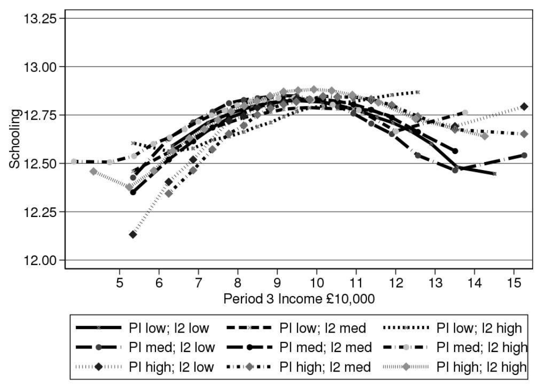


$\alpha_1 = 0.373\,(0.026)$ $\alpha_2 = -0.136\,(0.004)$

c) Household Income 12-17

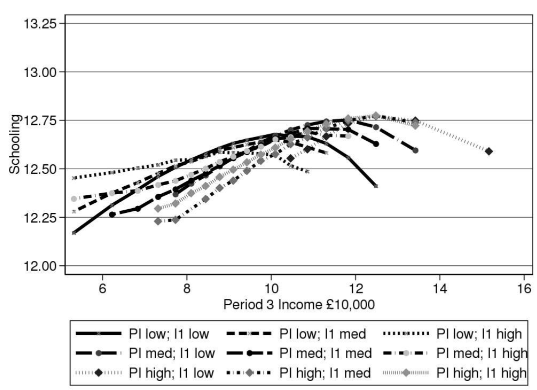


$\alpha_1 = 0.264\,(0.025)$ $\alpha_2 = 0.010\,(0.028)$

Note: Income in 2000 prices, £ 10,000s. Semiparametric regression controls and interpretation of $\alpha_1$ and $\alpha_2$ and associated standard errors is the same as in Figure 1. Low, medium, high levels of permanent income (PI), I1, I2, I3 indicate income at the 3rd, 5th and 7th decile. An adjustment of the scale of the x-axis and for 4b) the y-axis as described in Section 5.1.

Figure 5: Semi-parametric Estimates. Dependent Variable is Log Earnings at 30.

a) Household Income 6-11

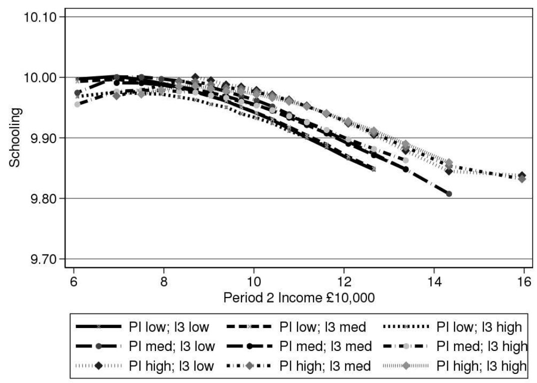


$\alpha_1 = -0.016\,(0.005)$ $\alpha_2 = -0.116\,(0.006)$

b) Household Income 12-17

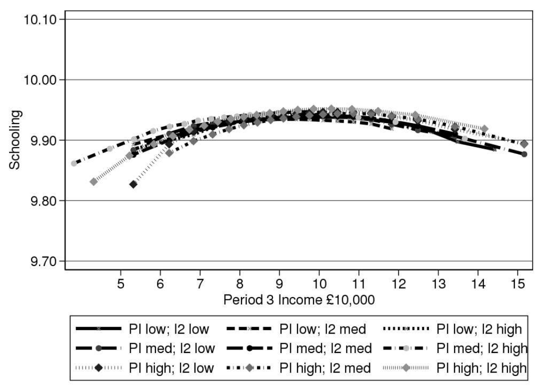


$\alpha_1 = 0.068\,(0.005)$ $\alpha_2 = -0.037\,(0.006)$

c) Household Income 12-17

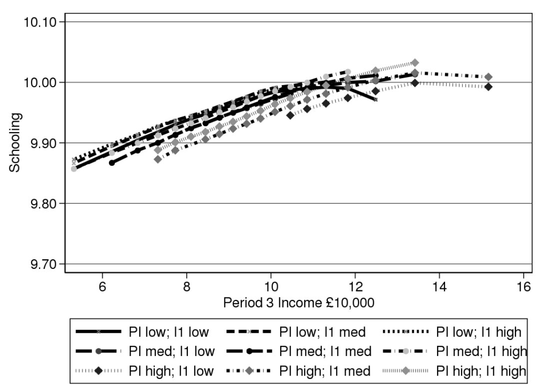


$\alpha_1 = 0.074\,(0.005)$ $\alpha_2 = 0.048\,(0.005)$

Note: Income in 2000 prices, £ 10,000s. Semiparametric regression controls and interpretation of $\alpha_1$ and $\alpha_2$ and associated standard errors is the same as in Figure 1. Low, medium, high levels of permanent income (PI), I1, I2, I3 indicate income at the 3rd, 5th and 7th decile. An adjustment of the scale of the x-axis and for 5b) the y-axis as described in Section 5.1.

Figure 6: Semi-parametric Estimates. Dependent Variable is IQ.

a) Household Income 6-11

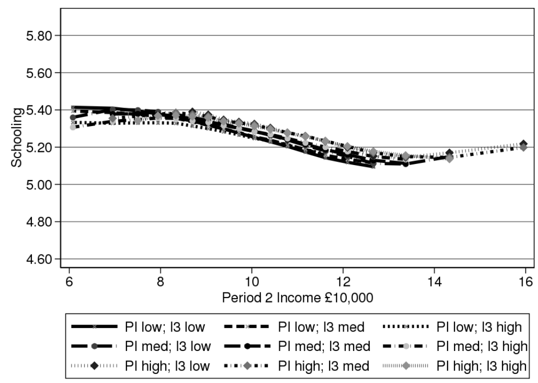


$\alpha_1 = -0.064\,(0.010)$ $\alpha_2 = -0.152\,(0.011)$

b) Household Income 12-17

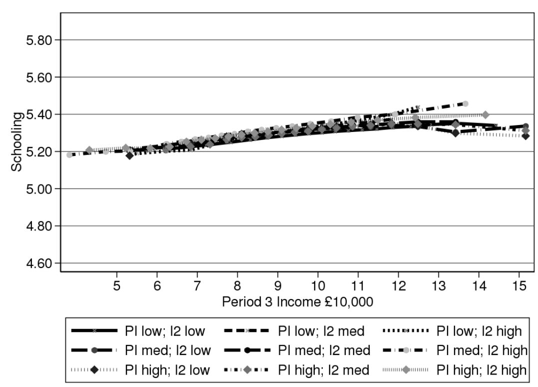


$\alpha_1 = 0.103\,(0.010)$ $\alpha_2 = 0.064\,(0.012)$

c) Household Income 12-17

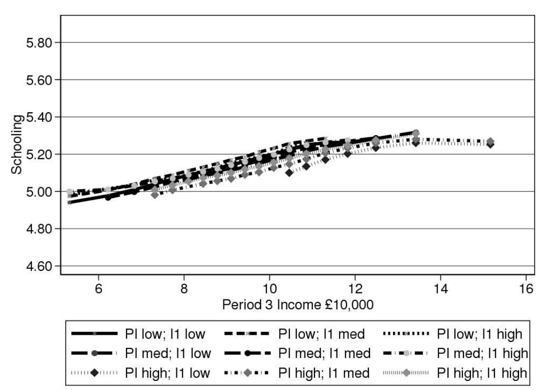


$\alpha_1 = 0.139\,(0.010)$ $\alpha_2 = 0.144\,(0.012)$

Note: Income in 2000 prices, £ 10,000s. Semiparametric regression controls and interpretation of $\alpha_1$ and $\alpha_2$ and associated standard errors is the same as in Figure 1. Low, medium, high levels of permanent income (PI), I1, I2, I3 indicate income at the 3rd, 5th and 7th decile. An adjustment of the scale of the x-axis and for 6b) the y-axis as described in Section 5.1.

Figure 7: Semi-parametric Estimates. Dependent Variable is Teen Pregnancy.

a) Household Income 6-11

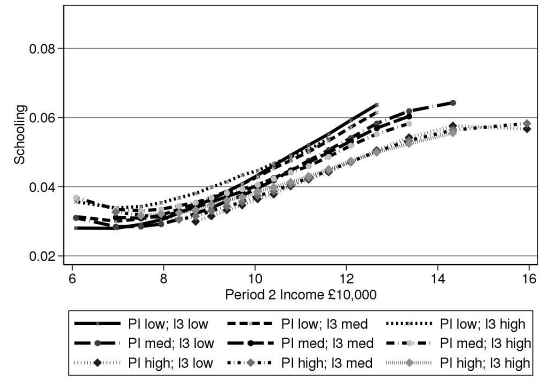


$\alpha_1 = 0.007\,(0.001)\ \alpha_2 = 0.021\,(0.002)$

b) Household Income 12-17

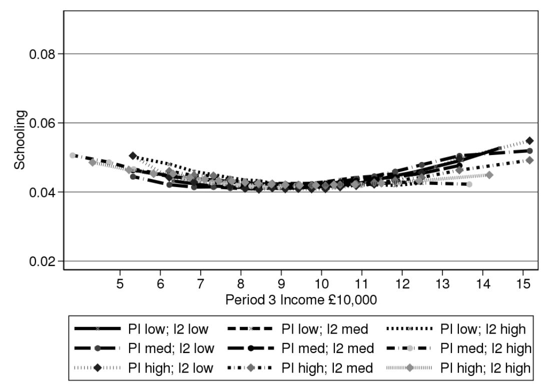


$\alpha_1 = -0.005\,(0.001)\ \alpha_2 = 0.006\,(0.001)$

c) Household Income 12-17

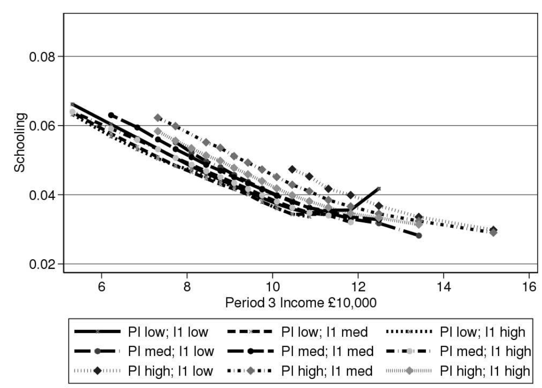


$\alpha_1 = -0.017\,(0.001)\ \alpha_2 = -0.011\,(0.001)$

Note: Income in 2000 prices, £ 10,000s. Semiparametric regression controls and interpretation of $\alpha_1$ and $\alpha_2$ and associated standard errors is the same as in Figure 1. Low, medium, high levels of permanent income (PI), I1, I2, I3 indicate income at the 3rd, 5th and 7th decile. An adjustment of the scale of the x-axis and for 7b) the y-axis as described in Section 5.1.

Table 1: Descriptive Statistics

| | N | Mean | Standard Deviation |
|---|---|---|---|
| Household Income | | | |
| Household Income Period 1, Age 0-5 | 499,442 | 10.72 | 3.30 |
| Household Income Period 2, Age 6-11 | 499,442 | 10.55 | 3.30 |
| Household Income Period 3, Age 12-17 | 499,442 | 9.95 | 4.24 |
| Household Permanent Income, Age 0-17 | 499,442 | 31.22 | 9.35 |
| Covariates | | | |
| Mother Years of Schooling | 499,442 | 11.15 | 2.71 |
| Father Years of Schooling | 499,442 | 11.47 | 3.01 |
| Mother Age at Birth | 499,442 | 26.25 | 5.02 |
| Father Age at Birth | 499,442 | 28.99 | 5.71 |
| Child Year of Birth | 499,442 | 1975.27 | 2.89 |
| Child Outcomes | | | |
| Years of Schooling | 499,442 | 12.74 | 2.41 |
| High School Dropout | 499,442 | 0.21 | 0.41 |
| College Attendance | 499,442 | 0.40 | 0.49 |
| Log Earnings age 30 | 295,380 | 9.97 | 0.79 |
| IQ (males only) | 238184 | 5.26 | 1.79 |
| Teenage Pregnancy (females only) | 238611 | 0.08 | 0.28 |

Note: Income and earnings variables in 2000 prices, £ 10,000s.